arXiv:cs/0303001v1 [cs.CG] 1 Mar 2003

When Crossings Count — Approximating the Minimum Spanning Tree*

Sariel Har-Peled† Piotr Indyk‡

February 1, 2008

Abstract

We present an $(1+\varepsilon)$-approximation algorithm for computing the minimum-spanning tree of points in a planar arrangement of lines, where the metric is the number of crossings between the spanning tree and the lines. The expected running time of the algorithm is near linear. We also show how to embed such a crossing metric of hyperplanes in d-dimensions, in subquadratic time, into high-dimensions so that the distances are preserved. As a result, we can deploy a large collection of subquadratic approximations algorithms [IM98, GIV01] for problems involving points with the crossing metric as a distance function. Applications include MST, matching, clustering, nearest-neighbor, and furthest-neighbor.

1 Introduction

Given a set of lines in the plane a natural measure of distances between any two points is the number of lines one has to cross to reach from one point to the other. This is a discrete distance measure that can be used to approximate the Euclidean distance and other distance measures. However, since this measure is defined by an arrangement of lines it is not locally defined and is thus computationally cumbersome. Finding the minimum spanning tree (MST) of a set of points, so that the number of intersections between the tree and the given set of lines is minimized, quantify how the set of points interact with the set of lines; see Figure 1. In fact, when the set of lines is the set of all possible lines, then this MST is the standard Euclidean MST [AF97] (here one minimizes the average number of edges of the MST crossed when picking a random line). Such an MST is related to a spanning tree of low stabbing number (STLSN) [Wel92, Aga91]. While the spanning tree of low stabbing number guarantee that any line intersects at most $O(\sqrt{n})$ edges of the spanning tree, the

*A preliminary version of the paper appeared in the *16th ACM Symposium of Computational Geometry*, 166–175, 2000.

†Department of Computer Science, University of Illinois, Urbana, IL, 61801, USA; sariel@cs.uiuc.edu; http://www.uiuc.edu/~sariel/

‡MIT Laboratory for Computer Science; 545 Technology Square, NE43-373; Cambridge, Massachusetts 02139-3594; indyk@theory.lcs.mit.edu

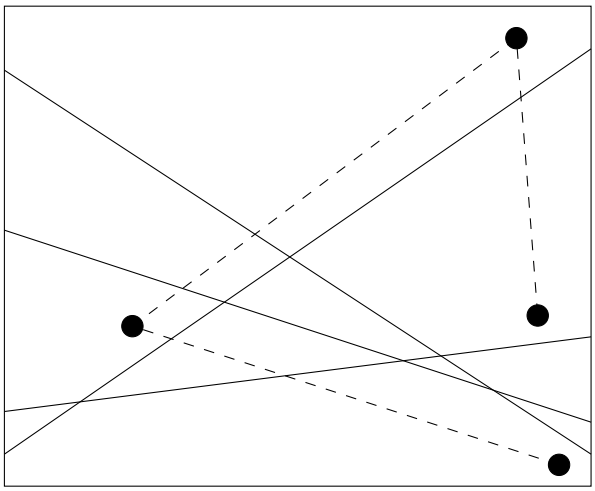

Figure 1: A set of lines and points, and the resulting crossing MST. Note that in this case the crossing MST is different from the Euclidean MST.

MST guarantees that the overall number of intersections between the tree and a given set of lines is minimized. Thus, if we have the set of lines in advance then the MST will have overall less intersections than the STLSN. The spanning tree of low-stabbing number was used in several applications, see for example [Aga91, MWW91]. In particular, having such an MST enables one: (i) to answer half-plane range queries in an efficient manner using a near linear space [GHS91], (ii) bound the complexity of the faces of the arrangement of lines that contain the points [HS01], and (iii) traverse between the points in an efficient way, so that the number of updates needed is minimized. (Imagine traversing among the points and maintaining the set of half-planes that contain the current point. Each time one crosses a line an update operation is performed.)

Computing the MST for the general case of arcs can be done in $O(n^2 \log n)$ time by performing wavefront propagation from each of the points (see Section 2). As for approximation algorithms, Har-Peled and Sharir [HS01] gave recently an approximation algorithm for the case of arcs, computing a Steiner tree in expected running time $w = O(\lambda_{t+2}(n+\mathcal{W}_{opt}) \log(n))$, where t is the maximum number of intersections between a pair of arcs, $\lambda_{t+2}(\cdot)$ is the maximum length of a Davenport-Schinzel sequence of order t, and $\mathcal{W}_{opt}$ is the weight of the optimal Steiner tree.[1] The algorithm outputs a tree of weight w (and thus gives roughly $O(\log n)$-approximation).

In this paper, we present two results:

- A near linear time $(1 + \varepsilon)$-approximation algorithm to the minimum-spanning tree under the crossing metric in the planar case.
- We show to embed the crossing metric among hyperplanes into a Hamming distance in high dimensions. As a result, we show how one can apply known subquadratic approximation algorithms for problems involving point-sets and hyperplanes in high dimensions (MST, clustering, matching, etc).

[1] It is easy to verify that if we have triangle inequality then the Steiner tree weight is at least half the weight of the MST.

The connection between the crossing metric, and points in high dimension follows by interpreting the input points as points in abstract VC-space [PA95] induced by the lines. Namely, we associate with each point in the plane, an n-dimensional binary vector, where i-th coordinate indicate on which side of the i-th line the point lies. In this way, we mapped our input points into points lying on the n-dimensional hypercube. The crossing metric is no more than the Hamming distance between the mapped points. We can now deploy the techniques of [IM98] to those mapped points, yielding an approximation algorithm for the MST problem. Bringing down the running time to be subquadratic requires some additional work.

Specifically, we show how to compute a mapping of the points into space of dimension $O(\log^7 n)$; this embedding can be computed in $\widetilde{O}(n^{4/3})$ time[2], for n points, so that we get a $(1+\varepsilon)$ gap property for a specified range of distances is preserved.

As a result, we can solve several approximation problems for this metric, among them is the MST problem. In fact, our near-linear approximate MST algorithm in the plane can be roughly viewed as an unraveling of the corresponding MST approximation algorithm in high dimensions. Similar bounds can be derived for $d > 2$ dimensions. See Section 4 for details.

The paper is organized as follows: In Section 2, we describe how one can compute the exact MST using wavefront propagation. In Section 3, we present the planar $(1+\varepsilon)$-approximation algorithm for the MST. Next, in Section 4, we describe the embedding into points in high dimension and demonstrate its usage for computing an approximate MST. Concluding remarks are given in Section 5.

2 Minimum Spanning Tree by Continuous Dijkstra

In this section, we present a simple algorithm for computing the crossing MST. It relies on a simple direct solution interpreted as a geometric algorithm. We also present a "weight sensitive" algorithm (Lemma 2.7) that computes portions of the MST in time proportional to its overall weight.

In the following, we assume that we are given a set L of lines and a set P of points in the plane. For simplicity, we assume $|P| = |L| = n$.

Definition 2.1 For a set L of lines, the *crossing metric* is defined to be the minimum number of lines of L that one has to cross as one moves between two prespecified points. Thus, for a pair of points $p, q \in \mathbb{R}^2$ the crossing distance between p and q, denoted by $\mathcal{D}_L(p,q)$, is the number of lines of L that intersects the segment pq. If L is a set of arcs, a similar crossing metric is defined, although the "shortest path" in this case is no longer necessarily a straight segment.

Definition 2.2 For a set L of lines, and a set P of points in the plane, let $\mathcal{T}_{opt}(P,L)$ denote a minimum spanning tree of P under the crossing metric induced by L, and let $\mathcal{W}_{opt}(P,L)$ denote the weight of $\mathcal{T}_{opt}(P,L)$.

[2]Here and in the rest of this paper $f(n) = \widetilde{O}(g(n))$ iff $f(n) = g(n)(1/\varepsilon)^{O(1)} \log^{O(1)} n$, and $f(n) = O_\varepsilon(g(n))$ iff $f(n) = O(g(n)/\varepsilon^{O(1)})$

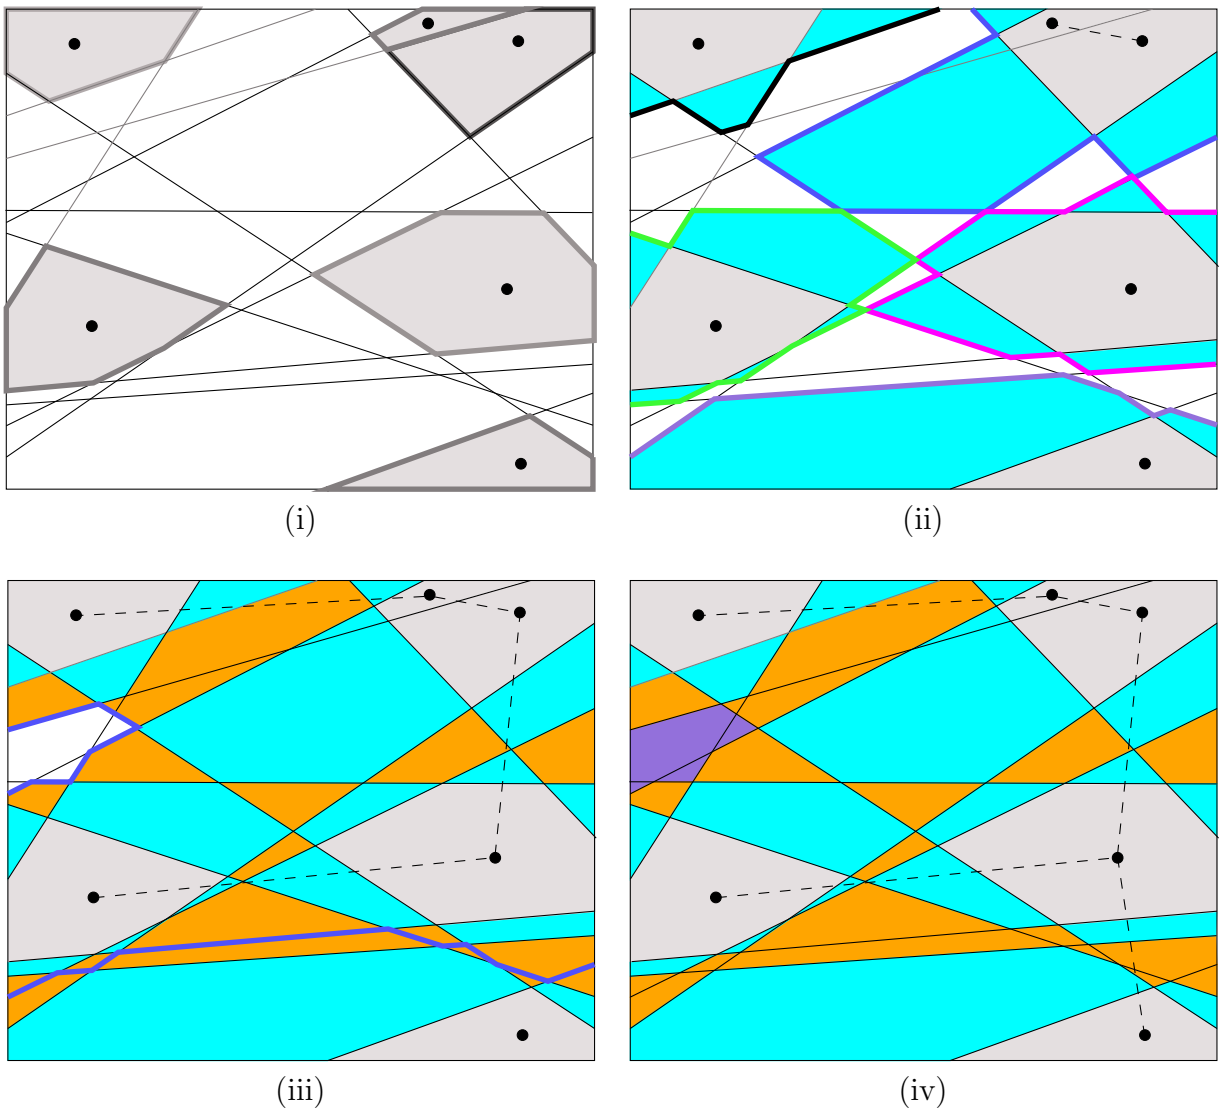

Figure 2: Computing the crossing MST by doing wavefront propagation. The thick lines denote the boundary of the current connected components of the spanning forest.

Let $\mathcal{A} = \mathcal{A}(L)$ denote the planar arrangement induced by the lines of L. Let $G_{adj} = G_{adj}(\mathcal{A})$ be the adjacency graph of $\mathcal{A}$; namely, each face of $\mathcal{A}$ is a vertex, and two vertices are connected if the two corresponding faces share an edge. Let V be the set of vertices of G_{adj} that corresponds to the faces of $\mathcal{A}$ that contains points of P. Clearly, the crossing MST of P in $\mathcal{A}$, corresponds to the MST of V in the graph G_{adj} (here, each edge has associated weight 1).

Computing the MST of V in G_{adj} can be done by performing a simultaneous flooding of G_{adj} from the vertices of V. Indeed, we compute in the i-th iteration all the vertices of G_{adj} that are in distance $\leq i$ from any vertex of V. This can be easily done using a modified BFS. In the beginning, the flood front is made out of n connected components. Every time two connected components of the flood front collide, we discovered a new edge of the MST. This edge connects the two vertices that induced the two parts of the wavefront that collided. This is a somewhat non-standard algorithm for computing the MST, but one can easily verify that it indeed computes the MST of V in G_{adj}.

This flooding algorithm has a natural geometric interpretation: Let $\mathcal{F}_{2i}$ denote the set of all faces of $\mathcal{A}$ that are in (crossing) distance at most i from any point of P. Clearly, $\mathcal{F}_0$ is the set of faces of $\mathcal{A}$ that contain points of P. The algorithm works in $n/2$ phases. We do a wavefront propagation in G_{adj}, starting from all the vertices that correspond to the marked faces (i.e., faces of $\mathcal{A}$ that contain points of P). In each iteration, we propagate the wavefront from the faces of $\mathcal{F}_{2i-2}$ into the faces of $\mathcal{F}_{2i}$. It is easy to verify that a connected component of the flood corresponds to a connected component of the wavefront of $\mathcal{F}_{2i}$. (Note, that two faces of $\mathcal{F}_{2i}$ might be adjacent but belong to different wavefronts as the wavefronts did not cross the separating edge yet and thus were not merged into a single wavefront.) The connected components are maintained implicitly by a union-find data-structure. In particular, during the i-th iteration of the wavefront propagation in G_{adj}, when two different connected components of the wavefront collide, it corresponds to two points of P with crossing distance equal to $2i-1$ or $2i$ from each other.

In particular, if there is an edge of the MST of weight $2i-1$ or $2i$ it would be discovered when the corresponding wavefronts collide. The i-th iteration of the wavefront propagation, corresponds to the detection of edges of weight $2i-1$ and $2i$ in the MST. For the MST applications, we first handle all relevant edges of weight $2i-1$, and later all such edges of weight $2i$. This requires a somewhat careful implementation, and we omit the the technical but straightforward details. See Figure 2.

Note, that the wavefront propagation can be done without constructing G_{adj} in advance, and one can compute parts of G_{adj} on the fly as needed (i.e., we need to compute only the parts of G_{adj} that are covered by the wavefront, or are about to be covered). Of course, in the worst case, the whole graph G_{adj} would be computed, which takes $O(n^2 \log n)$ time (this corresponds to computing the whole arrangement $\mathcal{A}(L)$).

Lemma 2.3 *Given a set L of n lines, and a set P of n points, a minimum spanning tree $\mathcal{T}_{opt}(P, L)$ of P under the crossing metric $\mathcal{D}_L$ can be computed in $O(n^2 \log n)$ time.*

Remark 2.4 In the algorithm of Lemma 2.3 we did not use the fact that L is a set of lines. The same algorithm will work for the case where L is a set of arcs. Since we do not have the triangle inequality in this case, the edges of the MST are no longer line segments, but rather

a Jordan arcs. (For example, imagine that the set L is a single segment and we would like to connect two points that are separated by this segment. This can be done with no crossing by going "around" this segment.)

To be able to generate parts of G_{adj} incrementally, as we perform the wavefront propagation, we need a way to compute the relevant portions of $\mathcal{A}(L)$ on the fly.

Theorem 2.5 ([HS01]) *Let L be a set of n lines, as above, and P a set of m points in the plane. Then one can compute, in expected $O((n+w+m)\,\alpha(n)\log n)$ time, a Steiner tree $\widehat{\mathcal{M}}$ of P, so that the expected weight of $\widehat{\mathcal{M}}$ is $O((n+w)\alpha(n)\log n)$, where $w = \mathcal{W}_{opt}(P,L)$ and $\alpha(n)$ is the inverse of the Ackermann function. Alternatively, one can compute the m faces that contain the points of P in the same time bound.*

Lemma 2.6 ([Aga91]) *There exists a Steiner tree $\mathcal{M}'$ of P, so that $\mathcal{W}_{opt}(P,L) = O(n\sqrt{n})$, and this is tight in the worst case (even for the case the arcs are lines).*

In the worst case, Theorem 2.5 is inferior to implicit point-location data-structures [AMS98] (which can perform the implicit point-location needed in roughly $O(n^{4/3})$ time for $m = n$), as implied by Lemma 2.6 (as the weight of the MST is $\Omega(n^{3/2})$ in the worst case, and this is the time to compute the relevant portions of the arrangement using the algorithm of Theorem 2.5). However, the running time of the algorithm of Theorem 2.5 is sensitive to the overall weight of the MST. This would be crucial for our algorithm.

Lemma 2.7 *Given a set L of n lines, a set P of n points, and a parameter i, one can compute, in expected $O(i(n+\mathcal{W}_{opt})\alpha^2(n)\log n)$ time, a minimum spanning forest of P under the crossing metric $\mathcal{D}_L$, that connects all the points of P in distance at most $\leq 2i$ from each other, where $\mathcal{W}_{opt} = \mathcal{W}_{opt}(P,L)$.*

Proof: The wavefront propagation on G_{adj} can be done using an implicit representation of the arrangement of $\mathcal{A}(L)$. Namely, we compute the set $\mathcal{F}_i$ of faces of $\mathcal{A}(L)$ in distance i from the points of P. Observe that the complexity of $\mathcal{F}_i$ is $O((n+\mathcal{W}_{opt})i\alpha(n/i))$. Indeed, the points of P can be connected by an arc $\gamma = \mathcal{T}_{opt}(P,L)$ having $O(\mathcal{W}_{opt})$ intersections with the lines of L, and let $\mathcal{A}'$ be the arrangement resulting from $\mathcal{A}$ by creating a tiny gate for each intersection of γ with the lines of L. The zone of γ in $\mathcal{A}(L)$ corresponds to a single face F of $\mathcal{A}'$, and the faces of $\mathcal{F}_i$ are contained in the set of faces in distance $\leq i$ from F. By [dBDS95], the complexity of this region is $O((n+\mathcal{W}_{opt})i\alpha(n/i))$ (this is a bound on the complexity of all the vertices in distance $\leq i$ from the face F.).

Clearly, the faces of $\mathcal{F}_i$ have a spanning tree of weight $O((n+\mathcal{W}_{opt})i\alpha(n/i))$, and so it can be computed in an online fashion in $O((n+\mathcal{W}_{opt})i\alpha^2(n)\log n)$ expected time, by Theorem 2.5. ∎

3 Approximation Algorithm for the Planar Case

The algorithm is depicted in Figure 3, Figure 4 and Figure 5. We next describe the algorithm and its analysis in more detail.

ALGORITHM `ApproxMST`(P, L, ε)

`Input:` A set of points P, a set of lines L, and an approximation parameter ε

`Output:` A spanning tree of P of weight $\leq (1+\varepsilon)\mathcal{W}_{opt}(P, L)$

begin

$M \leftarrow$ Approximate the weight of MST using the algorithm of Lemma Lemma A.6.

$l_0 \leftarrow \max\left(\frac{\varepsilon M}{c_5 n \alpha(n) \log^2 n}, 1\right)$

Set $F = (P, \emptyset)$ to be the an empty spanning forest of P.

`PropagateApproxWavefront`(P, L, l, F)

$i \leftarrow 1$

while F is not a single connected component **do**

$l_i \leftarrow l_{i-1} \cdot 2$

`PropagateApproxWavefront`(P, L, l, F)

$i \leftarrow i + 1$

end while

return F

end `ApproxMST`

Figure 3: Approximating the MST in the Plane

Lemma 2.7 provides us with an algorithm for approximating the MST in roughly quadratic time in the worst case. To get a near linear running time, we simulate the Dijkstra algorithm by performing the wavefront propagation in an approximate fashion.

Definition 3.1 A metric $\mathcal{D}'$ *ε-approximates* a metric $\mathcal{D}$, if for any $p, q, r, s \in P$ such that $\mathcal{D}'(p, q) \leq \mathcal{D}'(r, s)$ then $\mathcal{D}(p, q) \leq (1+\varepsilon)\mathcal{D}(r, s)$.

Definition 3.2 For a set F of segments in the plane, and a metric $\mathcal{D}$, let $\text{weight}_{\mathcal{D}}(F) = \sum_{e \in F} \mathcal{D}(e)$ denote the total weight of F under the metric D.

The proof of the following lemma is straightforward, and is included only for the sake of completeness.

Lemma 3.3 *Let the metric $\mathcal{D}'$ be an ε-approximation to the metric $\mathcal{D}$ over a point-set P. Let T' be an MST of P under $\mathcal{D}'$. Then,* $\text{weight}_{\mathcal{D}}(T') \leq (1+\varepsilon)\text{weight}_{\mathcal{D}}(T)$*, where T is the MST of P under $\mathcal{D}$, and* $\text{weight}(T)$ *is the total weight of the edges of T.*

Proof: Let $e'_1, \ldots, e'_{n-1}$ be the the edges of T' sorted by their weight $\mathcal{D}'(e'_1) \leq \ldots \leq \mathcal{D}'(e'_{n-1})$. Let $T_0 = T$, and let T_i be the tree resulting from removing the heaviest edge (according to $\mathcal{D}'$) from the cycle present in $T_{i-1} \cup \{e'_i\}$ (if e'_i is already in T_{i-1} we do nothing). Let e_i denote this removed edge. Clearly, $\mathcal{D}'(e'_i) \leq \mathcal{D}'(e_i)$ and, by definition, $\mathcal{D}(e'_i) \leq (1+\varepsilon)\mathcal{D}(e_i)$. Namely, we replaced an edge e_i by an edge e'_i which is heavier by a factor of $(1+\varepsilon)$. In the end of the process T_{n-1} is just T', and $\text{weight}_{\mathcal{D}}(T') \leq \sum_{i=1}^{n-1}(1+\varepsilon)\text{weight}_{\mathcal{D}}(e_i) \leq (1+\varepsilon)\text{weight}_{\mathcal{D}}(T)$. ∎

Lemma 3.3 suggest that if we can find a computationally cheaper approximate metric than $\mathcal{D}_L(\cdot, \cdot)$, then we can use it to compute the MST. A natural way to do that, is to

```
ALGORITHM PropagateWavefront( P, R, l, F )
   Input:  P - set of points
           R - set of lines
           l - propagation distance
           F - current spanning forest
   Output: An updated forest F with any pair of points of distance ≤ 2l in a
           single connected component
begin
   Initialize the data-structure D(R) of [HS01] for online point-location.
   Set W_0 to be the set of faces of A(R) that contains points of P.
         Use D(R) to compute those faces.
   for i = 1, ..., l do
      W_i ← Set of faces of A(R) of distance = i from points of P.
            Do wavefront propagation from W_{i-1}, and use D(R) to retrieve
            the faces of interest in A(R).
      if two different wavefronts collide then
         Add an edge connecting the two corresponding points to F
         Merge the corresponding connected components.
   end for
end PropagateWavefront
```

Figure 4: Doing the wavefront propagation

randomly sample a subset $R \subseteq L$, and use $\mathcal{D}_R(\cdot,\cdot)$ as the approximate metric. However, it is easy to verify that $\mathcal{D}_R$ is an ε-approximate metric to $\mathcal{D}_L$ only if $L = R$.

Definition 3.4 Let $\mathcal{D}', \mathcal{D}$ be two metrics, $\varepsilon > 0$, and l be prescribed parameters. The metric $\mathcal{D}'$ is an (ε, l)-*approximation* to $\mathcal{D}$, if for any $p, q, r, s \in P$, such that (i) $\mathcal{D}(p,q), \mathcal{D}(r,s) \geq l$, and (ii) $\mathcal{D}'(p,q) \leq \mathcal{D}'(r,s)$, we have $\mathcal{D}(p,q) \leq (1+\varepsilon)\mathcal{D}(r,s)$.

Namely, $\mathcal{D}'$ ε-approximates $\mathcal{D}$ for distances not smaller than l.

Definition 3.5 For l, ε, let $\nu(l,\varepsilon) = \max\bigl(128 c_{samp} \frac{\log n}{l\varepsilon^2}, 1\bigr)$, where c_{samp} is an appropriate constant. Let $\mathcal{RS}(L, l, \varepsilon)$ be a random subset of L generated by picking independently each line of L with probability $\nu(l,\varepsilon)$.

Let $\rho(l,\varepsilon) = \nu(l,\varepsilon) l = 128 c_{samp} \frac{\log n}{\varepsilon^2}$. The value $\rho(l,\varepsilon)$ is the expected crossing distance in $\mathcal{A}(\mathcal{RS}(L,l,\varepsilon))$ between two points $p, q \in P$ such that $\mathcal{D}_L(p,q) = l$.

Lemma 3.6 *Let L be a set of n lines in the plane, l a positive integer number, $\varepsilon > 0$, and let $R = \mathcal{RS}(L, l, \varepsilon)$ be a random subset of L.*

For any two points p, q of distance $\mathcal{D}_L(p,q) \geq l$ from each other we have

$$\mathcal{D}_L(p,q) \leq \frac{n}{r(1-\varepsilon/4)} \cdot \mathcal{D}_R(p,q) \leq (1+\varepsilon)\mathcal{D}_L(p,q),$$

with probability $\geq 1 - n^{-c_0}$.

Furthermore, $\mathcal{D}_R(\cdot,\cdot)$ is an (ε, l)-approximation to $\mathcal{D}_L(\cdot,\cdot)$ with high probability.

ALGORITHM `PropagateApproxWavefront`(P, L, l, F)

`Input:` P - set of points
L - set of lines
l - starting propagation distance
F - current spanning forest

`Output:` An updated forest F with any pair of points of distance $\leq 2l$ in a single connected component

begin

Compute a random sample R by choosing each line of L into the sample with probability $f(l) = 128c_6\frac{\log n}{l\varepsilon^2}$

`/* Approximate the wavefront propagation in` $A(L)$ `by doing it (exactly) in` $\mathcal{A}(R)$ `*/`

`PropagateWavefront`(P, R, $c_7 \log n/\varepsilon^2$, F)

`/*` c_7 `is an appropriate constant */`

end `PropagateApproxWavefront`

Figure 5: Doing the approximate wavefront propagation

Proof: Indeed, let $X_{pq} = D_R(p,q)$. We have,

$$\mu = E[X_{pq}] = \mathcal{D}_L(p,q) \cdot \nu(l,\varepsilon) \leq 128\mathcal{D}_L(p,q)c_{samp}\frac{\log n}{l\varepsilon^2} \geq \frac{128c_{samp}\log n}{\varepsilon^2}.$$

By Chernoff inequality [MR95, MPS98], we have that

$$\begin{aligned} P\left[|X_{pq}-\mu| > \frac{\varepsilon}{4}\mu\right] &\leq 2\left(\frac{e^{\varepsilon/4}}{\left(1+\frac{\varepsilon}{4}\right)^{1+\varepsilon/4}}\right)^{\mu} = 2\exp\Big(\mu\Big(\frac{\varepsilon}{4} - \Big(1+\frac{\varepsilon}{4}\Big)\log\Big(1+\frac{\varepsilon}{4}\Big)\Big)\Big) \\ &\leq 2\exp\left(\mu\left(\frac{\varepsilon}{4} - \Big(1+\frac{\varepsilon}{4}\Big)\left(\frac{\varepsilon}{4} - \frac{\varepsilon^2}{32}\right)\right)\right) \\ &\leq 2\exp\left(-\mu\frac{\varepsilon^2}{64}\right) \leq \exp\left(-\frac{128c_{samp}\log n}{\varepsilon^2}\cdot\frac{\varepsilon^2}{64}\right) \leq n^{-c_{samp}}, \end{aligned}$$

since $\log(1+x) \geq x - x^2/2$, for $0 \leq x \leq 1$. In particular, this implies that with high probability $\mu(1-\varepsilon/4) \leq X_{pq} \leq \mu(1+\varepsilon/4)$. Namely, with high probability we have

$$\begin{aligned} \mathcal{D}_L(p,q) &\leq \frac{X_{pq}}{\nu(l,\varepsilon)(1-\varepsilon/4)} \leq \frac{\nu(l,\varepsilon)(1+\varepsilon/4)}{\nu(l,\varepsilon)(1-\varepsilon/4)}\mathcal{D}_L(p,q) = \frac{1+\varepsilon/4}{1-\varepsilon/4}\mathcal{D}_L(p,q) \\ &\leq (1+\varepsilon)\mathcal{D}_L(p,q). \end{aligned}$$

Consider now four points p,q,r,s, such that $\mathcal{D}_L(p,q), \mathcal{D}_L(s,t) \geq l$ and $\mathcal{D}_R(p,q) \leq \mathcal{D}_R(r,s)$. By the above discussion, we have with high probability

$$\mathcal{D}_L(p,q)\cdot\nu(l,\varepsilon)(1-\varepsilon/4) \leq \mathcal{D}_R(p,q) \leq \mathcal{D}_R(r,s) \leq (1+\varepsilon)\mathcal{D}_L(r,s)\cdot\nu(l,\varepsilon)(1-\varepsilon/4).$$

Namely, $\mathcal{D}_L(p,q) \leq (1+\varepsilon)\mathcal{D}_L(r,s)$. Namely, $\mathcal{D}_R(\cdot,\cdot)$ is an (ε,l)-approximation to $\mathcal{D}_L(\cdot,\cdot)$ with probability $\geq 1 - \binom{n}{2}n^{-c_{samp}}$. ∎

Lemma 3.6 and Lemma 3.3 suggest that we compute the MST by computing an appropriate random sample R (by using a threshold l), and deploy the algorithms of Section 2 to compute the MST of P in $\mathcal{A}(R)$. Such an MST would be an approximate MST. There are two main problems with this approach: (i) For short distances (i.e., $l = 1$), just starting the wavefront propagation (i.e., Lemma 2.7) is prohibitively expensive (it roughly takes $O(\mathcal{W}_{opt}(P, L))$ time which might be $\Omega(n^{3/2})$), (ii) For long distances (i.e., $\geq i \cdot l$), the wavefront propagation becomes, again, prohibitly expensive (i.e. $\widetilde{O}(ni)$) by Lemma 2.7.

Corollary 3.7 *Let U be the total weight of all the edges of $\mathcal{T}$ having weigh less than $\varepsilon\mathcal{W}_{opt}(P, L)/(10n)$. Then $U \leq \varepsilon\mathcal{W}_{opt}(P, L)/10$.*

Lemma A.6 describes how we can approximate $\mathcal{W}_{opt}(P, L)$ to within a polylogarithmic factor using random sampling in near linear time. Since the algorithm of this lemma is very similar to the techniques used below, we defer its description to the appendix. Equipped with such approximation M, we know by Corollary 3.7 that we do not "care" about edges of the MST of length smaller than $l_0 = O(\varepsilon M/(n \operatorname{polylog}(n)))$. In particular, we can generate a random sample R_0 which provides an (ε, l_0)-approximation to $\mathcal{D}_L(\cdot, \cdot)$. Thus, we can approximate the MST by computing the MST of $\mathcal{T}_{opt}(P, R_0)$.

This, however, does not address the second problem. Indeed, computing the MST of $\mathcal{T}_{opt}(P, R_0)$ might still be too expensive, as the following lemma testifies.

Lemma 3.8 *Given a set L of n lines, a set P of n points, and parameters l, i, ε, U, such that $l = \Omega(\mathcal{W}_{opt}(P, L)/(nU))$ and let $R = \mathcal{RS}(L, l, \varepsilon)$ be a random sample of L. Then, one can compute, in expected $\widetilde{O}(iUn)$ time, a minimum spanning forest of P under the crossing metric $\mathcal{D}_R$, that connects all the points of P in distance at most $\leq 2i$ from each other.*

Proof: Let X denote the size of R. Clearly, The expected value of X is

$$E[X] = n\nu(l, \varepsilon) = 128nc_{samp}\frac{\log n}{l\varepsilon^2} = O\left(\frac{Un^2 \log n}{\varepsilon^2 \mathcal{W}_{opt}(P, L)}\right),$$

by Definition 3.5. Let $\gamma = \mathcal{T}_{opt}(\gamma, L)$. Let $Y = \operatorname{weight}(\gamma, R)$. Clearly,

$$E[Y] = \operatorname{weight}(P, R) = \mathcal{W}_{opt}(P, L)\nu(l, \varepsilon) = O\left(\frac{Un \log n}{\varepsilon^2}\right).$$

Namely, $E[\mathcal{W}_{opt}(P, R)] \leq E[Y] = O\big(\frac{Un \log n}{\varepsilon^2}\big)$. The running time bound now follows immediately by applying the algorithm of Lemma 2.7 to P and R. ∎

The algorithm of Lemma 3.8 first performs wavefront propagation for distances in $\mathcal{A}(R)$ which are smaller than $\rho(l, \varepsilon)$. For such distances $\mathcal{A}(R)$ *does not provide* reliable estimate (i.e., ordering) of the crossing distances between points. However, once the distances propagated exceed $\rho(l, \varepsilon)$, we know by Lemma 3.6 that the distances are now (ε, l)-approximated correctly. The main importance of the algorithm of Lemma 3.8 is that the algorithm has near linear running time for small values of U and i.

Using Lemma 3.8 together with Corollary 3.7 implies that we can compute a spanning forest for the "short" edges of $\mathcal{T}_{opt}(P, L)$ in near linear time.

Lemma 3.9 *Given a set P of n points in the plane, and a set L of n lines in the plane. One can compute a spanning forest F of P, such that the weight of F is $\leq \varepsilon \mathcal{W}_{opt}(P, L)/10$. Furthermore, every pair of points of P in distance $\Omega(\mathcal{W}_{opt}(P, L)\varepsilon/(n \log^3 n))$ belong to the same connected components of F. The running time of this algorithm is $\widetilde{O}(n)$.*

Proof: Using the algorithm of Lemma A.6, compute in $\widetilde{O}(n)$ time, a number M such that $\mathcal{W}_{opt}(P, L) \leq M = O(n\alpha(n) \log^2 n + \mathcal{W}_{opt}(P, L)\alpha(n) \log n)$. In particular, let

$$l_{short} = \frac{\varepsilon M}{c_1 n \log^3 n} \leq \frac{\varepsilon}{40n} \mathcal{W}_{opt}(P, L), \tag{1}$$

for c_1 large enough. On the other hand, $l_{short} = \Omega(\mathcal{W}_{opt}(P, L)/(Un))$, where $U = O((\log^3 n)/\varepsilon)$.

We now compute a spanning forest for P, using Lemma 3.8 with l_{short} and U as specified and $i = 2\rho(l, \varepsilon)$. The running time of this algorithm is

$$\widetilde{O}(iUn) = \widetilde{O}(\rho(l_{short}, \varepsilon)n) = \widetilde{O}\left(\frac{\log n}{\varepsilon^2} \cdot n\right) = \widetilde{O}(n)\,.$$

Clearly, F has at most n edges, and all the points of P in distance $\leq l_{short}$ are in the same connected component of F by Lemma 3.6.

Furthermore, for any edge pq of F, we have that with high probability $\mathcal{D}_L(p, q) \leq 2(1 + \varepsilon)l_{short} \leq 4l_{short}$ by Lemma 3.6. In particular, $\text{weight}(F, L) \leq 4nl_{short} \leq (\varepsilon/10)\mathcal{W}_{opt}(P, L)$. ■

Lemma 3.9 implies that we can compute a cheap spanning forest of P in near linear time that "captures" all the light edges of the MST. Next, we can compute the rest of the edges of the MST using Lemma 3.8 repeatedly.

Lemma 3.10 *Given a set P of n points in the plane, and a set L of n lines in the plane, a parameter $\varepsilon > 0$, and a spanning forest F of P, such that every pair of points of P in distance $\leq l$ belong to the same connected components of F, where $l = \Omega(\mathcal{W}_{opt}(P, L)\varepsilon/(n \log^3 n))$. Then, one can compute a spanning forest F' of P such that all the points of F in distance $\leq 2l$ belong to the same connected component of F'. The forest F' can be computed in $\widetilde{O}(n)$ expected time.*

Proof: We use the same algorithm of Lemma 3.9, with the modification that when calling to the algorithm of Lemma 3.8, we pass on F, such that the algorithm ignore generated edges that belong to the same connected component of F. It is again clear, that only edges of length between l and $2(1 + \varepsilon)l$ would be added to the spanning forest. The exact details of how to specify U and i are similar to Lemma 3.9, and are omitted. ■

Our algorithm for computing the MST works by using Lemma 3.9. This results in a spanning forest F_0 of the points of P, and a value l_{short} as specified by Equation (1). We now use Lemma 3.10 repeatedly $O(\log n)$ times, in the i-th iteration handling distances between $2^{i-1}l_{short}$ to $2 \cdot 2^i l_{short}(1 + \varepsilon)$, for $i = 1, \ldots, O(\log n))$, till we handle all distances $\leq n$. Namely, in the i-th iteration, we compute a spanning forest F_i of all points in distance $\leq 2^i l_{short}$ from each other using Lemma 3.10 using F_{i-1} as our "starting" spanning forest.

Clearly, the expected running time of the resulting algorithm is $\widetilde{O}(n)$. What is not clear, is that the resulting MST is indeed an ε-approximate MST.

Lemma 3.11 *With high probability, the tree T computed by the above algorithm is an ε-MST of P in $\mathcal{A}(L)$.*

Proof: All the edges generated by the algorithm of Lemma 3.9, in the first stage of the algorithm, have total weight $\leq (\varepsilon/20)\mathcal{W}_{opt}(P, L)$ with high probability.

Let $\mathcal{T}_{opt}(P, L)$ be the optimal spanning tree. If T is not an ε-approximate MST, then $\text{weight}_{\mathcal{D}_L}(T) > (1+\varepsilon)\text{weight}_{\mathcal{D}_L}(\mathcal{T}_{opt})$. In particular, there must be an edge of $\mathcal{T}_{opt}$ which its insertion into T would results in substantially lightly spanning tree. Formally, for an edge e, let $T(e)$ be the tree resulting from T by inserting e into T, and removing from T the heaviest (according to $\mathcal{D}_L$) edge on the new cycle that was created, and let $\text{out}(T, e)$ denote this "ejected" edge.

Arguing as in the proof of Lemma 3.3, it must be that there exists an edge $\phi = pq$ of $\mathcal{T}_{opt}$ such that

$$(1+\varepsilon)\mathcal{D}_L(\phi) < \mathcal{D}_L(\text{out}(T, \phi)),$$

and $\mathcal{D}_L(\phi) > \mathcal{W}_{opt}(P, L)/(20n)$.

Let i be the index such that $2^{i-1}l_{short} \leq \mathcal{D}_L(\phi) \leq 2^i l_{short}$. With high probability, we know that after the i-th iteration p and q are in the same connected component of F_i. Assume that p and q were not in the same connected component of F_{i-1} (the other case is easier and as such is omitted).

Let T'' be the spanning forest maintained by the algorithm just after p and q were present in the same connected component. With high probability, for any edge e'' of T'', we have $\mathcal{D}_L(e'') \leq (1+\varepsilon)\mathcal{D}_L(\phi)$, since the random sample R_i we used in the i-iteration is $(2^{i-1}l_{short}, \varepsilon)$-approximation to $\mathcal{D}_L$.

But then, it is not possible that the algorithm added $\text{out}(T, \phi)$ to the spanning tree T'', as all the edges on the cycle in $T'' \cup \{\phi\}$ are lighter than $(1+\varepsilon)\mathcal{D}_L(\phi)$. A contradiction. ■

We summarize our result:

Theorem 3.12 *Given a set P of n points in the plane, L a set of n lines, and $\varepsilon > 0$ a parameter. Then one can compute a spanning tree T of P, in $\widetilde{O}(n)$ expected time, such that $\text{weight}(T, L) \leq (1+\varepsilon)\mathcal{W}_{opt}(P, L)$. The result is correct with high probability.*

4 Approximation Algorithms for the Intersection Metric via Embeddings

Let $P = \{p_1, \ldots, p_n\}$ be a given set of n points, and $L = \{l_1, \ldots, l_m\}$ be a set of m lines, where $m = n^{O(1)}$. As mentioned earlier, the metric $\mathcal{D}_L$ is computationally cumbersome. One possible way to overcome this problem, is to embed this metric into a more convenient metric (while introducing a small distortion error).

In this section, we show a somewhat weaker result. We show how to embed the points of P into $O(\log^7 n)$-dimensional space in $\widetilde{O}(n + m + n^{2/3}m^{2/3})$ time, so that a specific distance gap in the crossing metric, is mapped to a corresponding gap in the target space.

We first observe that the crossing distance between two points p and q, can be computed by interpreting this distance as a Hamming distance on the hypercube in m dimensions induced by the lines. Namely, each line l contribute a coordinate — a point gets a '1' in

this coordinate if it is on one side of l, and a '0' if it is on the other side of l. Formally, let l^+ denote the open half-plane defined by a line l that contains the origin, and l^- denote the other open plane. For a point $p \in \mathbb{R}^2$, let $\vec{v}_L(p) = (b_1, \ldots, b_m)$ be a m-bit vector so that $b_i = 1$ **iff** $p \in l_i^+$. It is easy to verify that $\mathcal{D}_L(p,q) = d_H(\vec{v}_L(p), \vec{v}_L(q))$, where d_H is the Hamming distance.

Definition 4.1 Let $R \subseteq L$, let $f_R : \mathbb{R}^2 \to \mathbb{Z}$ be the mapping that maps a point p in the plane to its face ID in the arrangement $\mathcal{A}(R)$. Formally, we assign for each face in the arrangement $\mathcal{A}(R)$ a unique integer (say, and integer between 1 and $O(|R|^2)$). The mapping f_R maps a point p in the plane to the integer identifying the face that contains p. (Note, that is does not uniquely define $f_R(\cdot)$ as we did not specify how we assign the IDs to the faces.)

For a set $\mathcal{R} = (R_1, \ldots, R_\mu)$ of subsets of L, let $f_{\mathcal{R}} : \mathbb{R}^2 \to \mathbb{Z}^\mu$ be the mapping $f_{\mathcal{R}}(p) = (f_{R_1}(p), f_{R_2}(p), \ldots, f_{R_\mu}(p))$. For two points $p, q \in \mathbb{R}^2$, let $d_H(f_{\mathcal{R}}(p), f_{\mathcal{R}}(q))$ be the Hamming distance between $f(p)$ and $f_{\mathcal{R}}(q)$. Namely, this is the number of coordinates, where the two vectors $f_{\mathcal{R}}(p)$ and $f_{\mathcal{R}}(q)$ disagree.

One can view $f_{\mathcal{R}}$ as an embedding of the crossing metric $\mathcal{D}_L$ to the Hamming space $\mathbb{Z}^\mu$.

Lemma 4.2 *Given a set P of n points in the plane, a set L of lines in the plane, a parameter $\varepsilon > 0$ and a parameter r. One can compute a set $\mathcal{R}$ of μ subsets of L, such that for the embedding $f_{\mathcal{R}} : \mathbb{R}^2 \to \mathbb{Z}^\mu$, we have that, with high probability, for any $p, q \in P$ it holds:*

- *If $\mathcal{D}_L(p,q) \leq r$, then $d_H(f(p), f(q)) \leq M$,*
- *If $\mathcal{D}_L(p,q) \geq (1+\varepsilon)r$ then $d_H(f(p), f(q)) \geq (1+\varepsilon)(1 - a/\log n)M$,*

where M and a are appropriate constants and $\mu = O(\log^4 n)$.

Proof: For sake of simplicity of exposition, we assume that $m/r \geq \log n$, where $m = |L|$. If this is not correct, we can add "fictitious" lines to L that have all the points of P on one side of them. If we pick such a line to a set of $\mathcal{R}$, we can ignore it when we compute the face IDs.

For a parameter α to be specified shortly, let $k = \alpha m/r$, R be a sample of k lines out of L (performed with replacement), and let p, q be two points of P. Let $\rho = \mathcal{D}_L(p,q)/n$. The probability that p, q will be in two different faces of $\mathcal{A}(R)$ is

$$U(\rho) = 1 - (1-\rho)^k,$$

as this is the probability that not all the lines will miss the segment connecting p and q.

Our target is to approximate the value of $U(\rho)$ so we could decide whether p, q are close or far. Indeed, if $U(\rho) \geq U((1+\varepsilon)r/m)$ then $\mathcal{D}_L(p,q) \geq (1+\varepsilon)r$, and if $U(\rho) \leq U(r/m)$ then $\mathcal{D}_L(p,q) \leq r$.

To do so, we generate a set of subsets $\mathcal{R} = (R_1, \ldots, R_\mu)$, by random sampling as described above, where μ would be specified shortly. Now we consider the quality of the distance approximation provided by the embedding[3]. Let $X(p,q)$ denote the random variable which is

[3] A similar analysis (in the context of Hamming spaces) appeared already in [Ind00]; in our case, however, we have to put more care into the analysis, since we want ε and ε' to be very close.

the number of arrangements of $\mathcal{A}(R_1), \ldots, \mathcal{A}(R_\mu)$ that have p, q in different faces. Note, that $X(p,q)$ is equal to the Hamming distance between $f_{\mathcal{R}}(p)$ and $f_{\mathcal{R}}(q)$, and it thus the distance between the images of p and q in the new space. Clearly, as μ tends to infinity, $X(p,q)/\mu$ tends to $U(\rho)$. Using Chernoff inequality, we can quantify the quality of approximation provided by μ. Specifically, let $z = U(r/m)$ and $Z = U((1+\varepsilon)r/m)$; in the following we will make sure that $Z < 1/2$. Then, from the Chernoff bound [MR95, MPS98] it follows that for any $\alpha > 0$ if $\mu = C\frac{\log n}{z\alpha^2}$ for some constant C, then with high probability:

- if $\mathcal{D}_L(p,q) \leq r$ then $X(p,q)/\mu \leq z(1+\alpha)$
- if $\mathcal{D}_L(p,q) \geq r(1+\varepsilon)$ then $X(p,q)/\mu \geq Z(1-\alpha)$

Therefore, the mapping $f_{\mathcal{R}}$ converts the distance gap $r : (1+\varepsilon)r$ into the gap $z(1+\alpha)\mu : Z(1-\alpha)\mu$. We next fine tune k (the size of each sample) so that the resulting gap will be as large as possible. (Intuitively, the larger the target gap is, the easier it is to detect it in later stages.) Therefore, in the following we focus on finding k such that the ratio

$$\Delta = \frac{Z(1-\alpha)\mu}{z(1+\alpha)\mu}$$

is as large as possible. To this end, we observe that

$$\begin{aligned} z &= U\Big(\frac{r}{m}\Big) = 1 - \Big(1 - \frac{r}{m}\Big)^k \leq 1 - e^{-rk/m}\left(1 - \frac{(rk/m)^2}{k}\right) = 1 - e^{-\alpha}\left(1 - \frac{\alpha^2}{k}\right) \\ &\leq 1 - e^{-\alpha}\big(1-\alpha^2\big) \leq \alpha^2 + (1-e^{-\alpha})\big(1-\alpha^2\big) \leq \alpha^2 + \alpha\big(1-\alpha^2\big) \leq \alpha(1+\alpha) \end{aligned}$$

since $\left(1 - \frac{t}{n}\right)^n \geq e^{-t}\left(1 - \frac{t^2}{n}\right)$ [MR95], $k = \frac{\alpha m}{r}$, and $x \geq 1 - e^{-x}$. Furthermore,

$$\begin{aligned} Z &= U((1+\varepsilon)r/m) = 1 - (1 - (1+\varepsilon)r/m)^k \geq 1 - e^{-(1+\varepsilon)rk/m} = 1 - e^{-(1+\varepsilon)\alpha} \\ &\geq (1+\varepsilon)\alpha - ((1+\varepsilon)\alpha)^2 \geq (1+\varepsilon)\alpha(1 - (1+\varepsilon)\alpha) \end{aligned}$$

since $(1 - t/n)^n \leq e^{-t}$ [MR95] and $1 - e^{-x} \geq x - x^2/2 \geq x - x^2$.

Therefore

$$\frac{Z}{z} \geq \frac{(1+\varepsilon)\alpha(1-(1+\varepsilon)\alpha)}{\alpha(1+\alpha)} \geq (1+\varepsilon)(1-(1+\varepsilon)\alpha)(1-\alpha) \geq (1+\varepsilon)(1-(2+\varepsilon)\alpha),$$

since $1/(1+x) \geq (1-x)$. Thus, if we set α to be $1/\log n$, then the distance gap becomes (at least)

$$\Delta = \frac{Z(1-\alpha)\mu}{z(1+\alpha)\mu} \geq (1+\varepsilon)(1-(2+\varepsilon)\alpha)(1-\alpha)^2 \geq (1+\varepsilon)\left(1 - \frac{a}{\log n}\right),$$

where a is an appropriate constant. Also, note that the resulting value of z is

$$z = 1 - (1 - r/m)^k \geq 1 - e^{-p_0 k} = 1 - e^{-\alpha} \geq \alpha - \alpha^2/2 = \Omega(1/\log n)$$

and $\mu = (C\log n)/(z\alpha^2) = C\log^2 n/\alpha^2 = O(\log^4 n)$. Finally, since $m/r \geq \log n$, we have that $k = \alpha(m/r) = (1/\log n)(m/r) \geq 1$ (i.e., the sample size k is at least 1). ■

Lemma 4.3 *Given a set P of n points, and a set L of m lines, one can compute the function $f_{\mathcal{R}}(\cdot)$, of Lemma 4.2, for all the points of P in $\widetilde{O}((m^{2/3}n^{2/3}+m+n))$ expected time.*

Proof: We have to compute for each point of P the face that contains it in each of the arrangements $\mathcal{A}(R_1),\ldots,\mathcal{A}(R_\mu)$, where $\mu = O(\log^4 n)$. Or alternatively, compute all the faces of $\mathcal{A}(R_1),\ldots,\mathcal{A}(R_\mu)$ that contains points of P. For a single arrangement A_i this can be done in
$O(m^{2/3}n^{2/3}\log^{2/3}(m/\sqrt{n}) + (m+n)\log m)$ expected time [AMS98]. Since there are μ coordinates (i.e., arrangements), the result follows. ∎

Thus, we showed how to embed $\mathcal{D}_L$ into μ-dimensional Hamming space Σ^μ in $\widetilde{O}(n + m + n^{2/3}m^{2/3})$ time, mapping a $(1+\varepsilon)$ gap between close and far points into a gap of size $(1+\varepsilon)(1-O(1)/\log n)$, where $\mu = O(\log^4 n)$ and $\Sigma \subseteq \mathbb{Z}$ is the set of face labels we use (i.e., $|\Sigma| = O(m^2)$. By using standard embedding techniques (e.g. see [KOR00]) we can embed the Hamming space Σ^μ into $\{0,1\}^D$ with $D = O(\mu \log|\Sigma| \log^2 n) = O(\log^6 n \log m)$, preserving the gap up to another factor $(1-O(1)/\log n)$. This gives an embedding of $\mathcal{D}_L$ into $D = O(\mu \log m \log^2 n)$-dimensional binary Hamming cube, with error $(1-O(1)/\log n)$. Thus it is sufficient for us to maintain c-nearest neighbor in $\{0,1\}^D$ where $c = (1+\varepsilon)(1-O(1)/\log n)$, which takes $\widetilde{O}(n^{1/c}) = \widetilde{O}(n^{1/(1+\varepsilon/2)})$ time per operation [IM98].

We conclude:

Theorem 4.4 *By performing a $\widetilde{O}(n+m+n^{2/3}m^{2/3})$-time preprocessing, one can reduce the problem of maintaining dynamic $(1+\varepsilon)$-approximate nearest neighbor for any n-point crossing metric over m lines, to the problem of maintaining dynamic $(1+\varepsilon)(1-O(1)/\log n)$-approximate nearest neighbor in Hamming space with $O(\log^7 n)$ dimensions (assuming $m = n^{O(1)}$). The latter can be solved in $\widetilde{O}(n^{1/(1+\varepsilon/2)})$ time per operation.*

4.1 Embedding of the Crossing Metric over $\mathbb{R}^d$

In this Section, we extend the methods from the previous section to the crossing metric defined by $(d-1)$-dimensional hyperplanes in $\mathbb{R}^d$, for any fixed $d \geq 2$. To this end, it is sufficient to design an efficient procedure, which given a set of n points $P = p_1,\ldots,p_n$ and a set of m hyperplanes $\mathcal{H} = \{H_1,\ldots,H_m\}$, assigns a symbol $\sigma_i \in \Sigma \subset \mathbb{Z}$ to each p_i in such a way that $\sigma_i \neq \sigma_j$ iff there exists H_k which separates p_i from p_j. Unfortunately, the idea from the previous section does not give subquadratic time algorithm for $d > 2$, since even in $d = 3$ the complexity of n cells in an arrangement formed by n planes could be $\Omega(n^2)$. Fortunately, for our purpose, we do not need to compute the actual cells containing p_is. Rather, it is just sufficient to find the *labels* for those cells, or more specifically, a function $h : P \to \Sigma$ such that $h(p) = h(q)$ iff p and q belong to the same arrangement cell.

Abusing notations, we denote by $H_k(p)$ the function returning 1 if p lies on one side of H_k and zero otherwise. We use the following hashing function

$$h(x) = \left(\sum_i a_i H_i(x)\right),$$

where $a_1 \ldots a_m$ are independent and identically distributed random variables with uniform distribution over $\{0,\ldots,n^c\}$, where c is a constant to be specified shortly. Note, that if

$p, q \in \mathbb{R}^d$ lie in two different full-dimensional faces of $\mathcal{A}(\mathcal{H})$, then, as noted above, there must be a hyperplane $H_k \in \mathcal{H}$, so that $H_k(p) \neq H_k(q)$, and say that $H_k(p) = 1$. That is, $h(p) = h'(p) + a_k$ and $h(q) = h'(q)$, where $h'(x) = \sum_{i \neq k} a_i H_i(x)$. Since the a_i were picked independently, it follows that $h(p) = h(q)$ only if $h'(p) - h'(q) = a_k$. But the probability of that to happen is $1/n^c$. We conclude, that the probability of two points belonging to two different faces to be mapped to the same value by $h(\cdot)$ is $1/n^c$. Thus, since we have $O(n^2)$ pairs of points to consider in our algorithm, it follows that the probability of the hashing to fail is n^{2-c} which can be made to be arbitrarily small by picking c to be large enough.

Namely, we associate a weight a_i with each half-space induced by a hyperplane H_i. For each point p_j, we compute the total weight of all the half-spaces that contain it, and all the points having the same total weight are associated with the same label. Computing the weight of a point p_j falls into the class of problems known as intersection-searching [Aga97]. In particular, one can construct a data-structure in $O(m^{1+\delta})$ time, so that one can answer intersection-searching queries in $O((n/m^{1/d}) \log^{d+1} n)$ time, where $\delta > 0$ is arbitrarily small constant. As the algorithm needs to perform a linear number of such queries, we set $m = n^{2d/(d+1)}$. Thus, the algorithm computes the required labels in $O(n^{2d/(d+1)+\delta})$ time. We conclude:

Theorem 4.5 *By performing a $O(n^{2d/(d+1)+\delta})$-time preprocessing, where $\delta > 0$ is arbitrary constant, one can reduce the problem of maintaining dynamic $(1+\varepsilon)$-approximate nearest neighbor for any n-point crossing metric over n hyperplanes in $\mathbb{R}^d$, to the problem of maintaining dynamic $(1+\varepsilon)(1 - O(1)/\log n)$-approximate nearest neighbor in Hamming space with $O(\log^7 n)$ dimensions.*

Remark 4.6 Note, that the constants in the bounds of Theorem 4.5 depend exponentially (or worse) on the dimension d.

Remark 4.7 As indicated in the introduction, having such a embedding, enable one to use a large collection of subquadratic approximation algorithms for the intersection metric, including dynamic amortized $\widetilde{O}(n^{4/3} + n^{1+1/c})$-time (for $d = 2$) c-approximation algorithms for bichromatic closest pair [Epp95] and $\widetilde{O}(n^{4/3} + n^{1+1/c})$-time algorithms for: c-approximate diameter and discrete minimum enclosing ball [GIV01], $O(c)$-approximate facility location and bottleneck matching [GIV01]. Similar (i.e., subquadratic time) results hold for any $d > 2$.

4.2 Computing an MST Using the Embedding

We next describe how to use the embedding described in the previous two sections, for getting an $(1+\varepsilon)$-approximation algorithm for the MST under crossing metric. Note that everything described in this section is well known [IM98], and we provide it only for the sake of completeness. Also, the resulting algorithm is slower in the planar case than the algorithm of Section 3.

Computing the minimum spanning tree under the intersection metric, using the Kruskal's algorithm, boils down to maintaining the bichromatic nearest-neighbor pair (under the intersection metric) between two sets $P_1, P_2 \subseteq P$, under insertions and deletions. A consequence of Eppstein result [Epp95] is the following:

Theorem 4.8 ([Epp95]) *Given a dynamic data-structure for nearest-neighbor queries, where each insertion / deletion / query operation takes $T(n)$ time, then one can compute the MST in $O(nT(n)\log^2 n)$ time.*

It is easy to verify that if we get a $(1+\varepsilon)$-approximation to the MST if we use an $(1+\varepsilon)$-approximate dynamic nearest-neighbor data-structure (Eppstein, personal communication, 1999).

Namely, we need a data-structure that support dynamic approximation nearest-neighbor queries. After applying the embedding described above, we use the ε'-PLEB data-structure of [IM98] to maintain a $(1+\varepsilon')$-approximate nearest neighbor in the embedded space. Specifically, we construct an ε-PLEB in the embedded points. In this way, we obtain an ε-PLEB for our original points (i.e., we embedded a gap to a gap, so that a close point in the embedded space, corresponds to a close point in the crossing metric) data-structure that for a query p return us a point of $q \in P$ so that $\mathcal{D}_L(p,q) \leq (1+\varepsilon)r$, if there exits a point $q^* \in P$ so that $\mathcal{D}_L(p,q^*) \leq r$.

Thus, by constructing $\log_{1+\varepsilon} n$ such data-structures, we can use binary search on those data-structures to find and $(1+\varepsilon)$-approximate nearest neighbor to a query point. Namely, this data-structure can be used to answer approximate nearest neighbor queries for the intersection metric. For the whole scheme to work, we need those data-structures to be dynamic; i.e., support insertions and deletions of points. Fortunately, the only part of the algorithm that needs to be dynamic is the second stage that uses the data-structure of [IM98] which is dynamic.

We conclude:

Theorem 4.9 *Given a set P of n points in the plane, and a set L of n lines, one can compute in $\widetilde{O}\left(n^{4/3} + n^{1+1/(1+\varepsilon)}\right)$ time, a spanning tree of P of weight $\leq (1+\varepsilon)\mathcal{W}_{opt}(P,L)$. The result returned by the algorithm is correct with high probability. For $d > 2$ dimensions, such an MST can be approximated in $\widetilde{O}\left(n^{2d/(d+1)+\delta} + n^{1+1/(1+\varepsilon)}\right)$ time, where $\delta > 0$ is an arbitrary constant.*

5 Conclusions

We presented the first $(1+\varepsilon)$-algorithm for approximating the minimum spanning tree under the crossing metric in the plane. We also presented a subquadratic time approximation algorithms for a variety of other problems, obtained by embedding the crossing metric into higher dimensional space. The techniques used in our paper seems to be new to low-dimension computational geometry, and we believe that they might be useful for other problems in computational geometry.

There are several interesting open problems for further research:

- Can the result be extended to other cases: segments or arcs instead of lines?
- Can a similar approximation algorithm be found for the case of minimum weight triangulation under the crossing metric?

Acknowledgments

The authors wish to thank Pankaj Agarwal, Boris Aronov and Micha Sharir for helpful discussions concerning the problems studied in this paper and related problems.

References

[AF97] B. Aronov and S. Fortune. Average-case ray shooting and minimum weight triangulations. In *Proc. 13th Annu. ACM Sympos. Comput. Geom.*, pages 203–211, 1997.

[Aga91] P. K. Agarwal. *Intersection and Decomposition Algorithms for Planar Arrangements*. Cambridge University Press, New York, NY, 1991.

[Aga97] P. K. Agarwal. Range searching. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 31, pages 575–598. CRC Press LLC, Boca Raton, FL, 1997.

[AMS98] P. K. Agarwal, J. Matoušek, and O. Schwarzkopf. Computing many faces in arrangements of lines and segments. *SIAM J. Comput.*, 27(2):491–505, 1998.

[dBDS95] M. de Berg, K. Dobrindt, and O. Schwarzkopf. On lazy randomized incremental construction. *Discrete Comput. Geom.*, 14:261–286, 1995.

[Epp95] D. Eppstein. Dynamic Euclidean minimum spanning trees and extrema of binary functions. *Discrete Comput. Geom.*, 13:111–122, 1995.

[GHS91] L. J. Guibas, J. Hershberger, and J. Snoeyink. Compact interval trees: a data structure for convex hulls. *Internat. J. Comput. Geom. Appl.*, 1(1):1–22, 1991.

[GIV01] A. Goel, P. Indyk, and K. R. Varadarajan. Reductions among high dimensional proximity problems. In *Proc. 12th ACM-SIAM Sympos. Discrete Algorithms*, pages 769–778, 2001.

[HS01] S. Har-Peled and M. Sharir. Online point location in planar arrangements and its applications. *Discrete Comput. Geom.*, 26:19–40, 2001.

[IM98] P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proc. 30th Annu. ACM Sympos. Theory Comput.*, pages 604–613, 1998.

[Ind00] P. Indyk. Dimensionality reduction techniques for proximity problems. In *Proc. 10th ACM-SIAM Sympos. Discrete Algorithms*, pages 371–378, 2000.

[KOR00] E. Kushilevitz, R. Ostrovsky, and Y. Rabani. Efficient search for approximate nearest neighbor in high dimensional spaces. *SIAM J. Comput.*, 2(30):457–474, 2000.

[MPS98] E. W. Mayr, H. J. Promel, and A. Steger, editors. *Lectures on Proof Verification and Approximation Algorithms*, volume 1367. Springer-Verlag, 1998.

[MR95] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, 1995.

[MWW91] J. Matoušek, E. Welzl, and L. Wernisch. Discrepancy and ε-approximations for bounded VC-dimension. In *Proc. 32nd Annu. IEEE Sympos. Found. Comput. Sci.*, pages 424–430, 1991.

[PA95] J. Pach and P. K. Agarwal. *Combinatorial Geometry*. John Wiley & Sons, New York, NY, 1995.

[Wel92] E. Welzl. On spanning trees with low crossing numbers. In *Data Structures and Efficient Algorithms, Final Report on the DFG Special Joint Initiative*, volume 594 of *Lecture Notes Comput. Sci.*, pages 233–249. Springer-Verlag, 1992.

A A Rough Approximation to the Weight of the MST in Near Linear Time

In this appendix, we show how to approximate the weight of the minimum spanning tree up to roughly a factor of $O(\alpha(n)\log n)$ if its weight is at least linear. In Section 3, we presented a near linear time algorithm for $(1+\varepsilon)$-approximation for the minimum spanning tree, that relies on this approximation algorithm.

Underlining the approximation algorithm, is the observation that an MST for a random sample of the lines of L provides a rough approximation to the weight of the MST of L. If the weight of the MST of the sample is near linear, we can approximate it up to a $O(\alpha(n)\log n)$, using the following algorithm.

Lemma A.1 *Given a set R of r lines, P a set of n points, and W a prescribed parameter, one can decide whether $\mathcal{W}_{opt}(P,R)$ is large; namely, $\mathcal{W}_{opt}(P,R) = \Omega((r+n+W)\alpha(n)\log n)$. The algorithm takes $O((r+n+W)\alpha(n)\log^2 n)$ expected time. Furthermore, if $\mathcal{W}_{opt}(P,R) \leq W$, the algorithm will report that its weight is large with probability at most n^{-c}, where c is an appropriate constant.*

Proof: Use the algorithm of Theorem 2.5 and execute it $O(\log n)$ times on P and R. If the running time of the i-th execution of the algorithm exceeds $\Omega((r+n+W)\alpha(n)\log n)$ abort it, and move on to the next execution. If $\mathcal{W}_{opt}(P,R) \leq W$, then the algorithm of [HS01] provides a spanning tree of expected weight $O((r+n+W)\alpha(n)\log n)$ with the same bound on the expected running time. Thus, if in $O(\log n)$ executions the algorithm returns always that $\mathcal{W}_{opt}$ is large, we can conclude that with probability $\geq 1-n^{-c}$ the weight of $\mathcal{W}_{opt}(P,R)$ is not $\leq W$. ∎

Lemma A.1 shows that we can approximate the weight of the MST in near linear time if its weight is near linear. However, if it is heavier, we will use random sampling to keep the running time under control.

Let $R \subseteq L$ be a random sample of lines out of L, where each line is picked independently with probability r/n. Clearly, the probability of an intersection point u (between a connected set γ and a line of L), to be present in $\mathcal{A}(R)$ is r/n (this is the probability that the line of L passing through u will be chosen to be in the random sample).

Definition A.2 For a curve γ, and a set of lines L, let $\mathrm{weight}(\gamma, L)$ denote the *weight* of γ in the arrangement $\mathcal{A}(L)$. This is the number of intersections of γ with the lines of L.

Lemma A.3 *Let R be a sample of lines of L (chosen as described above), then with high probability:*

$$\mathcal{W}_{opt}(P, L) \leq \frac{n}{r}(c_0 n \log n + 2\mathcal{W}_{opt}(P, R)),$$

and with probability ≥ 0.9 we have $\frac{n}{r} \cdot \frac{\mathcal{W}_{opt}(P,R)}{10} \leq \mathcal{W}_{opt}(P, L)$, where c_0 is an appropriately large constant.

Proof: Let $\mathcal{T}^L_{opt} = \mathcal{T}_{opt}(P, L)$, and let $W_R = \mathrm{weight}(\mathcal{T}^L_{opt}, R)$ be the weight of $\mathcal{T}^L_{opt}$ under the crossing metric of R. Clearly, $E[W_R] = \mathcal{W}_{opt}(P, L)\frac{r}{n}$. Thus, we know that with probability ≥ 0.9 we have $W_R \leq 10\mathcal{W}_{opt}(P, L)\frac{r}{n}$ (by Markov inequality), and with probability ≥ 0.9, we have that $\mathcal{W}^R_{opt} = \mathcal{W}_{opt}(P, R) \leq W_R \leq 10\mathcal{W}_{opt}(P, L)\frac{r}{n}$.

Let $p, q \in P$ be two points, and let X_{pq} be the distance between p, q in the arrangement $\mathcal{A}(R)$. If the distance between p, q is large, that is $U = \mathcal{D}_L(p, q) \geq c_0(n/r)\log n$ (where c_0 is a large enough constant), then one can show using Chernoff inequality, that with high probability, we have:

$$\frac{U}{2} \leq X_{pq}\frac{n}{r} \leq 2U.$$

On the other hand, by the above argument, each edge $e = pq$ of $\mathcal{T}^R_{opt} = \mathcal{T}_{opt}(P, R)$ either intersects at most $c_0(n/r)\log n$ lines of L, or alternatively, the number of lines of L intersected by e is smaller than $2(n/r)X_e$, where X_e is the number of lines of R that e intersects. Thus, with high probability, we have

$$\begin{aligned}
\mathcal{W}_{opt}(P, L) &\leq \mathrm{weight}(\mathcal{T}^R_{opt}, L) = \sum_{e=pq \in \mathcal{T}^R_{opt}} \mathcal{D}_L(p, q) \leq \sum_{e \in \mathcal{T}^R_{opt}} \left(c_0 \frac{n}{r} \log n + 2X_e \frac{n}{r} \right) \\
&= c_0 \frac{n^2 \log n}{r} + \mathcal{W}_{opt}(P, R)\frac{2n}{r}.
\end{aligned}$$

■

Remark A.4 We can make both probabilities in Lemma A.3 large by repeating the experiment $O(\log n)$ times, and picking the smallest $W(P, R)$ computed. With high probability, we have

$$\frac{n}{r} \cdot \frac{\mathcal{W}_{opt}(P, R)}{10} \leq \mathcal{W}_{opt}(P, L) \leq \frac{n}{r}(c_0 n \log n + 2\mathcal{W}_{opt}(P, R)).$$

In particular, if $\mathcal{W}_{opt}(P, R) > c_0 n \log n$, we get that $3\mathcal{W}_{opt}(P, R)\frac{n}{r}$ is a constant factor approximation to $\mathcal{W}_{opt}(P, L)$.

Lemma A.5 *Let r be a prescribed parameter, and $\mathcal{W}_{opt} = \mathcal{W}_{opt}(P, L)$. Then, an algorithm can decide whether*

- *$\mathcal{W}_{opt}$ is small - namely $\mathcal{W}_{opt} \leq \frac{10c_0 n^2 \log n}{r}$.*
- *$\mathcal{W}_{opt}$ is large - $\mathcal{W}_{opt} = \Omega(\frac{n^2}{r}\alpha(n)\log^2 n)$.*
- *$\mathcal{W}_{opt}$ is in between. Any of the two above answers are valid.*

The algorithm takes $O(n\alpha(n)\log^4 n)$ time, and returns a correct result with high probability.

Proof: We pick $m = O(\log n)$ samples $R_1, \ldots, R_m$ by picking each line with probability r/n into the sample. For each sample, we check whether $\mathcal{W}_{opt}(P, R_i) \leq 10c_0 n \log n$, using the algorithm of Lemma A.1. This will require $O(n\alpha(n)\log^3(n))$ time for each sample, and $O(n\alpha(n)\log^4(n))$ overall.

If the algorithm of Lemma A.1 returned *not large* for any sample R, we know that $\mathcal{W}_{opt}(P, R) = O(n\alpha(n)\log^2 n)$. And by Lemma A.3, we know that $\mathcal{W}_{opt}(P, L) = O\Big(\frac{n^2\alpha(n)\log^2 n}{r}\Big)$ with high probability. ∎

Now, we can perform a binary search to approximate the weight of $\mathcal{W}_{opt}(P, L)$.

Lemma A.6 *One can compute in $O(n\alpha(n)\log^5 n)$ time a value M, so that*

$$\mathcal{W}_{opt}(P, L) \leq M = O(n\alpha(n)\log^2 n + \mathcal{W}_{opt}(P, L)\alpha(n)\log n).$$

Proof: Use Lemma A.5, set $r_0 = n$. In the i-th iteration check whether $\mathcal{W}_{opt} = \Omega\Big(\frac{n^2}{r_i}\alpha(n)\log^2 n\Big)$, by using the algorithm of Lemma A.5. If it is, we set $r_{i+1} = r_i/2$, and repeat the process. We stop as soon as this check fails. Then, we know that with high probability

$$\frac{10c_0 n^2 \log n}{r_{i-1}} \leq \mathcal{W}_{opt}(P, L) = O\left(\frac{n^2\alpha(n)\log^2 n}{r_i}\right) = M,$$

implying that M is the required approximation. ∎

Remark A.7 Note, that if algorithm of Lemma A.6 stops after the first iteration, then $\mathcal{W}_{opt} = O(n\alpha(n)\log^2 n)$. In such a case the approximation we get is much worse then logarithmic. However, this is to some extent the easiest case: Without any sampling we get a spanning tree of near linear (or sub linear) weight.